Quantum Model Averaging

Christopher Ferrie
Center for Quantum Information and Control, University of New Mexico, Albuquerque, New Mexico, 87131-0001, USA
(Dated: May 27, 2014)

Standard tomographic analyses ignore model uncertainty. It is assumed that a given model generated the data and the task is to estimate the quantum state, or a subset of parameters within that model. Here we apply a model averaging technique to mitigate the risk of overconfident estimates of model parameters in two examples: (1) selecting the rank of the state in tomography and (2) selecting the model for the fidelity decay curve in randomized benchmarking.

arXiv:1405.6329v1 [quant-ph] 24 May 2014

Contents

I. Introduction 1

II. General problem and common methods 2
A. Problem setup 2
B. Akaike Information Criterion (AIC) 3
C. Bayesian Information Criterion (BIC) 3

III. Bayesian model selection and averaging 4
A. Model average estimate 4
B. Sequential Monte Carlo 5

IV. Examples 6
A. Rank selection 6
B. Randomized benchmarking 7

V. Conclusion and discussion 8

Acknowledgments 9

References 10

I. INTRODUCTION

Parameter estimation is an integral part of physics. Accurate estimates of physical parameters in quantum mechanical models allows for precision quantum control [1, 2] which enables practical goals such as quantum computation [3] and quantum metrology [4] which in turn can provide probes of fundamental physics such as gravity wave detection [5].

Quantum parameter estimation shares many similarities with its classical counterpart. But there are many subtle and peculiar differences. Even in single parameter estimation, quantum metrology [4, 7] shows that we can obtain advantages from quantum resources such as squeezed states of light [9] and entanglement [10]. With such rich structure, many new subtleties [11], considerations [12–14] and generalities [15–17] can arise, including entirely new approaches to estimation [18–21] and verification [22–24].

At the other end of the parameter spectrum the quantum state of a physical system is our most complete description of it. Thus estimation of quantum states [6] might be seen as the most ambiguous form of estimation. There are many approaches to the general problem [25–28], some which specialize for computational efficiency [29–32] and those which go beyond to estimation of regions [33–38].

All of the above mentioned results assume a model. That is, it was taken as given that a particular parametric distribution generated the data. But what if this assumption is not correct? This question has garnered recent interest and classical approaches to *model selection* have been used in a variety of experimental [43, 44] and theoretical [39–42] works. Here we supplement these results with a new approach to parameter estimation: *model averaging*. The technique shares many similarities with model selection—in fact, model selection is a crucial component of model averaging, but goes beyond it. Although model selection adds an additional layer of security to overconfident

estimation, selecting models can be itself a red herring, for the most probable model might only be slightly more probable than others.

The approach considered here combines Bayesian parameter estimation with Bayesian model selection, such that the final estimate of the parameters is the best value of the parameters within each model, averaged over the probability assigned to each model. We will show that such an approach can reduce the error incurred by first selecting a model—which has some probability of being incorrect—then selecting parameters within that model. In fact, the numerical experiments presented here show that model average estimation always does better than the estimates from incorrect models and in some scenarios can perform better than even estimates from the correct model. This is due to the additional hedging afforded by considering multiple models, all of which carry *some* information.

The paper is organized as follows. In Sec. II we outline the problem and review the model selection techniques used so far. In Sec. III, we present the full Bayesian approach to model selection and define the *model average estimate* of parameters. Sec. IV presents two distinct examples where the model average estimate provides an advantage for parameter estimation. We conclude with a discussion in Sec. V.

II. GENERAL PROBLEM AND COMMON METHODS

In Sec. II A we overview the problem. In Secs. II B and II C we review the Akaike information criterion (AIC) and the Bayesian information criterion (BIC) which are by far the most commonly used approaches to model selection (and the only used thus far for quantum estimation). These sections are included for reference and completeness.

A. Problem setup

Let us begin with the base problem of *parameter estimation*. A physical model prescribes the probabilities for the outcomes of experiments: $\Pr(D|\boldsymbol{x};C,M)$. Here D is some hypothetical or observed data, $\boldsymbol{x}$ is a set of real numbers in $\mathbb{R}^d$, where d is the *dimension* of the model, C is the experimental context[1] and M is our model. For example, we could have the model M be that of qubit which is parameterized by a Bloch vector $\boldsymbol{x} = (x, y, z)$. The experimental context could be a measurement basis, say that of σ_x. Then, the quantum mechanical model prescribes $\Pr(\pm|\boldsymbol{x};C,M) = (1 \pm x)/2$. In quantum mechanics this is called the *Born rule* and in statistics, the *likelihood function*.

Going from parameters to the probability of data is a deductive process—the model gives us numerical values of $\Pr(D|\boldsymbol{x};C,M)$. The experiment, on the other hand, gives us a particular data set D—but what we really want is $\boldsymbol{x}$. This is an example of an *inverse problem*. What the Bayesian solution provides is $\Pr(\boldsymbol{x}|D;C,M)$, the *distribution* of $\boldsymbol{x}$ given the data. Although we lack certainty about $\boldsymbol{x}$, we can accurately (read: quantitatively, mathematically) describe our state of knowledge of $\boldsymbol{x}$ given we have seen the data D.

The formal path forward is through Bayes' rule:

$$\Pr(\boldsymbol{x}|D;C,M) = \frac{\Pr(D|\boldsymbol{x};C,M)\Pr(\boldsymbol{x}|C,M)}{\Pr(D|C,M)}. \tag{1}$$

From a chronological point of view relative to D, we begin with the *prior* $\Pr(\boldsymbol{x}|C,M)$, which encodes the information we have about $\boldsymbol{x}$ prior to learning about D. We weight the prior by the likelihood function and normalize by the *marginal likelihood*[2]:

$$\Pr(D|C,M) = \mathbb{E}_{\boldsymbol{x}|C,M}[\Pr(D|\boldsymbol{x};C,M)]. \tag{2}$$

The distribution produced as result of Bayes rule, $\Pr(\boldsymbol{x}|D;C,M)$, is called the *posterior* and represents our knowledge of $\boldsymbol{x}$ after the data has been observed.

At this point we could call the problem solved. This, however, assumes the model is correct. If the model is suspect, then we have the meta–problem of determining the best model.

---

[1] Experimental context, C, here means any additional information necessary to provide a well defined function $\Pr(D|\boldsymbol{x};C,M)$.

[2] The notation $\mathbb{E}_x[f(x)]$ means take the expectation of $f(x)$ with respect to the distribution of x (the distribution itself is implicit).

B. Akaike Information Criterion (AIC)

The most commonly used model selection[3] technique is the Akaike information criterion (AIC), which arises as follows. First, we suppose there is some true model, $M = T$, giving a distribution $\Pr(D'|T;C)$. We quantity the discrepancy between this model and our candidate model via the *Kullback-Leibler divergence*:

$$\mathrm{KL}(T\|M) = \mathbb{E}_{D'|T;C}\left[\log\frac{\Pr(D'|T,C)}{\Pr(D'|\boldsymbol{x};M,C)}\right] = \mathbb{E}_{D'|T;C}\left[\log\Pr(D'|T,C)\right] - \mathbb{E}_{D'|T;C}\left[\log\Pr(D'|\boldsymbol{x};M,C)\right], \tag{3}$$

for some appropriate choice of parameters $\boldsymbol{x}$. The "best" model, then, is the one that minimizes $\mathrm{KL}(T\|M)$, which is equivalent to maximizing the second term in (3) since the first term only depends on the true model. However, we do not know the true model and we do not know the best set of parameters within our candidate model. The latter problem is naturally address by collecting data D and producing an estimate of the parameters $\hat{\boldsymbol{x}}(D)$, then averaging over possible data sets such that the quantity of interest becomes

$$\mathbb{E}_{D|T;C}\left[\mathbb{E}_{D'|T;C}\left[\log\Pr(D'|\hat{\boldsymbol{x}}(D);M,C)\right]\right]. \tag{4}$$

Akaike showed that, independent of the true model (and under some regularity conditions), an unbiased estimator of this quantity is

$$\mathrm{AIC}(M) = \max_{\boldsymbol{x}}\log\Pr(D|\boldsymbol{x};C,M) - d, \tag{5}$$

where d is the number of dimension of the model (the number of free parameters). The preferred model is the one with largest value of $\mathrm{AIC}(M)$. The simple linear penalization with dimension makes it is clear how models with more parameters are penalized.

C. Bayesian Information Criterion (BIC)

The Bayesian approach is more general. Being so, it is less obvious how it might penalize complex models. Here we show how an asymptotic approximation leads to a form similar to the AIC. First, we write the marginal likelihood (2) as the integral expectation

$$\Pr(D|M;C) = \int d\boldsymbol{x}\,\Pr(\boldsymbol{x}|M;C)\Pr(D|\boldsymbol{x};C,M). \tag{6}$$

We will approximate this integral using Laplace's method. To this end, consider the Taylor expansion of the log of the likelihood function about its peak (note that for brevity, we have dropped the context C and model M from the conditionals):

$$l(\boldsymbol{x}) := \log\Pr(D|\boldsymbol{x}) \approx l(\boldsymbol{x}_0) + \nabla^{\mathrm{T}}l(\boldsymbol{x}_0)(\boldsymbol{x}-\boldsymbol{x}_0) + \frac{1}{2}(\boldsymbol{x}-\boldsymbol{x}_0)^{\mathrm{T}}\nabla\nabla^{\mathrm{T}}l(\boldsymbol{x_0})(\boldsymbol{x}-\boldsymbol{x_0}). \tag{7}$$

If we assume the number of measurements $N \to \infty$, the law of large numbers gives us

$$\nabla\nabla^{\mathrm{T}}l(\boldsymbol{x_0}) = \sum_{j=1}^{N}\nabla\nabla^{\mathrm{T}}l_j(\boldsymbol{x_0}), \tag{8}$$

$$\to \sum_{j=1}^{N}\mathbb{E}[\nabla\nabla^{\mathrm{T}}l_j(\boldsymbol{x_0})], \tag{9}$$

$$= -\sum_{j=1}^{N}I_j(\boldsymbol{x}_0), \tag{10}$$

$$= -NI(\boldsymbol{x}_0), \tag{11}$$

[3] See [46] for an overview of model selection and the various techniques mentioned here.

where I_j is the Fisher information of the j'th measurement and I is the arithmetic average of these values. Then, the integral (6) becomes

$$\Pr(D|M;C) = \int d\boldsymbol{x} \exp(l(\boldsymbol{x})) \Pr(\boldsymbol{x}), \tag{12}$$

$$= \exp(l(\boldsymbol{x}_0)) \Pr(\boldsymbol{x}_0) \left(\frac{2\pi}{N}\right)^{\frac{d}{2}} |I(\boldsymbol{x}_0)|^{-\frac{1}{2}} \left(1 + O\left(\frac{1}{N}\right)\right). \tag{13}$$

Now we take the logarithm to obtain

$$\log \Pr(D|M;C) \approx l(\boldsymbol{x}_0) - \frac{d}{2}\log N + \frac{d}{2}\log 2\pi + \log \Pr(\boldsymbol{x_0}) + \log |I(\boldsymbol{x}_0)|^{-\frac{1}{2}}. \tag{14}$$

If we ignore the terms not changing with N, we have a new quantity

$$BIC(M) = \max_{\boldsymbol{x}} \log \Pr(D|\boldsymbol{x};C,M) - \frac{d}{2}\log N, \tag{15}$$

which is the well-known Bayesian information criterion or BIC. Notice the striking similarity to the AIC (5). Being nearly equivalent, the BIC is often considered in addition to the AIC. Next, we will consider the full solution, which will allow us to obtain more accurate estimates of parameters averaged over the competing models.

Recent proposals have used the AIC/BIC on both simulated [39–42] and experimental data [43, 44]. In [42], however, the authors caution its unattended use. The argument against AIC for quantum states, for example, is simple. The AIC is derived from a metric which measures their closeness of model in their *predictive* probability—a certainly well-motivated measure. However, such a measure is only useful if all future measurements will be the same as those used to perform the data analysis. That is, one can measure copies of a quantum system in some fixed set of bases, estimate the state, then use that estimate to predict the outcome of a measurement in a new basis. Thus, in quantum theory, one ought to consider a measure on predictive distribution as maximized over all possible future measurements. As suggested in [42], such a measure might well be the quantum relative entropy, for example. Here we avoid these problems by considering the full Bayesian solution, which we describe next.

III. BAYESIAN MODEL SELECTION AND AVERAGING

A. Model average estimate

Within the Bayesian framework, the model selection approach is no different than for parameter estimation. Rather than focus on the distribution $\Pr(\boldsymbol{x}|D;C,M)$, we first consider $\Pr(M|D;C)$. Using Bayes rule we have

$$\Pr(M|D;C) = \frac{\Pr(D|M;C)\Pr(M|C)}{\Pr(D|C)}. \tag{16}$$

Often, practitioners of Bayesian methods go one step further and compare two models—say, M_1 and M_2—by taking the ratio of these posteriors,

$$\frac{\Pr(M_1|D;C)}{\Pr(M_2|D;C)} = \frac{\Pr(D|M_1;C)\Pr(M_1|C)}{\Pr(D|M_2;C)\Pr(M_2|C)}, \tag{17}$$

noticing the normalization factor cancels. This quantity is called the *posterior odds ratio* and first considered by Jeffreys [47]. Clearly, if the posterior odds ratio is larger than 1, we favor M_1. The last fraction is called the *prior odds ratio* and the unbiased choice favoring neither model is set this term equal to 1. This leaves us with

$$\frac{\Pr(M_1|D;C)}{\Pr(M_2|D;C)} = \frac{\Pr(D|M_1;C)}{\Pr(D|M_2;C)}, \tag{18}$$

which is called the *Bayes factor* [48]. Each quantity in the ratio is marginal likelihood (2) of its respective model.

For a discrete number of hypothetical models $\{M_k\}$, and assuming one model must be chosen, the optimal strategy is to compute the marginal likelihood if each model and select the one with the highest value. Model selection of this type has been used in quantum mechanical problems of Hamiltonian finding [49] and estimating error channels from syndrome measurements [50].

Here we will investigate the idea of using a meta-model, which is an average over those in $\{M_k\}$. Assume that we are interested in some subset of parameters $\boldsymbol{y}$ common to all models. Given we have taken data D and computed each marginal likelihood $\Pr(M_k|D;C)$, we define the model average estimate (MAE) as

$$\hat{\boldsymbol{y}}_{\text{MAE}}(D;C) = \mathbb{E}_{M_k|D;C}[\mathbb{E}_{\boldsymbol{y}|D;C,M_k}[\boldsymbol{y}]]. \tag{19}$$

In words, this is the average (over models) of the average (over parameters within models). Variants of this approach are referred to as Bayesian model averaging [45].

Before moving to our examples, a few comments are in order. First, it is not necessary that the models are "nested" in the sense that we can order them into supersets. The only requirement is that the parameters of interest are included in each model. The other parameters in each model we might call *nuisance parameters*. Part of the appeal of the Bayesian approach is that these parameters are automatically dealt with and we can focus on those parameters which are of immediate interest. Of course, the nuisance parameters can be inferred as well.

Let us wax philosophical for a moment. What is being proposed is to select a meta-model, an average over many different physical models. This may seem awkward for physical theories—after all, there is only one true model, right? Not in our view. Models are human constructs, those Platonic ideals which describe another world, a world we were clever enough to find through our mastery of mathematics and abstraction. Here, we have dropped the idea that the point is to find the capital-T-Truth. Rather, we measure our understanding of Nature through our ability to predict and control its behavior. By averaging physical models, we can show that this idea has merit.

B. Sequential Monte Carlo

In practice, the Bayesian update rule and the expectations required in the equations above are analytically and computationally intractable since they involve complicated integrals over multidimensional parameter spaces which may include solutions to equations of motions which are themselves intractable. To perform the calculations we turn to Monte Carlo techniques. Our numerical algorithm fits within the subclass of Monte Carlo methods called *sequential Monte Carlo* (SMC) or *particle filtering* [51].

The SMC procedure prescribes that we approximate the probability distribution by a weighted sum of Dirac delta-functions,

$$\Pr(\boldsymbol{x}) \approx \sum_{j=1}^{n} w_j \delta(\boldsymbol{x} - \boldsymbol{x}_j), \tag{20}$$

where the weights at each step are iteratively calculated from the previous step via

$$w_j \mapsto \Pr(D|\boldsymbol{x}_j) w_j, \tag{21}$$

followed by a normalization step. The elements of the set $\{\boldsymbol{x}_j\}_{j=0}^{n}$ are called *particles*. Here, $n = |\{\boldsymbol{x}_i\}|$ is the number of particles and controls the accuracy of the approximation. Like all Monte Carlo algorithms, the SMC algorithm approximates expectation values, such that

$$\mathbb{E}_{\boldsymbol{x}}[f(\boldsymbol{x})] \approx \sum_{j=1}^{n} w_j f(\boldsymbol{x}_j). \tag{22}$$

In other words, sequential Monte Carlo allows us to efficiently compute multidimensional integrals with respect to the measure defined by the probability distribution.

The resultant posterior probability provides a full specification of our knowledge. However, in most applications, it is sufficient—and certainly more efficient—to summarize this distribution. In our context, the optimal single parameter vector to report is the mean of the posterior distribution

$$\hat{\boldsymbol{x}} = \mathbb{E}_{\boldsymbol{x}}[\boldsymbol{x}] = \sum_{j=1}^{n} w_j \boldsymbol{x}_j. \tag{23}$$

The SMC approximation can also provide efficient calculation and description of *regions* [38, 52]. For our purpose, we also require the SMC approximation to give an accurate and efficient esimate of the marginal likelihood Eq. (6),

which we need to calculate Bayes' rule at the level of models Eq. (16). Via the SMC approximation, the integral expectation in the definition of the marginal likelihood, Eq. (6), is

$$\Pr(D) \approx \sum_{j=1}^{n} w_j \Pr(D|\boldsymbol{x}_j). \tag{24}$$

It is not immediate obvious but is easy to see in hindsight that this is exactly the normalization that must be computed already in the SMC algorithm after the weight update, Eq. (21), is applied. By storing this value, we can apply Bayes rule at the meta-level of models Eq. (16).

An iterative numerical algorithm such as SMC requires care to ensure stability. Conditions for stability of the algorithm and the specifications of an implementation have been detailed elsewhere [52, 53]. The SMC algorithm has now been used in many quantum mechanical parameter estimation problems [38, 49, 52–55] and a software implementation (the one used here) is available as a Python package [56, 57].

IV. EXAMPLES

A. Rank selection

We consider first an example similar to Guta, Kypraios and Dryden [44]: n qubits subjected to random Pauli measurements where the models under consideration are those of differing rank. Each model will be denoted M_r where r is the rank of the unknown quantum state so that the dimension of model M_r is $d = 2^{n+1}r$. We generate unknown quantum states with fixed rank r as follows [58]. Begin with a matrix $X \in \mathbb{C}^{2^n \times r}$ where each component x_{ij} is chosen independently according to $\Re(x_{ij}) \sim \mathcal{N}(0,1)$ and $\Im(x_{ij}) \sim \mathcal{N}(0,1)$—standard Normal distributions. Then define the rank r density operator

$$\rho = \frac{XX^\dagger}{\mathrm{Tr}(XX^\dagger)}. \tag{25}$$

If $r = 2^n$, this construction is equivalent to a Hilbert-Schmidt random density matrix. The model parameters will be the vectorization of the matrix X: $\boldsymbol{x} = \mathrm{vec}(X)$.

We label the single qubit Pauli operators $\{\sigma_0, \sigma_1, \sigma_2, \sigma_3\}$ and the multi-qubit Paulis by

$$\sigma_k = \sigma_{k_1} \otimes \sigma_{k_2} \otimes \cdots \otimes \sigma_{k_n}, \tag{26}$$

where $k = k_1 + 4k_2 + 4^2k_3 + \cdots + 4^{n-1}k_n$. Since each Pauli is idempotent, $\sigma_k^2 = \mathbb{1}$, each individual measurement has 2 possible outcomes which we label $d \in \{+1, -1\}$ for the $+1$ and -1 eigenvalues. Then the likelihood function of a single measurement can be related to the expectation value via: $\Pr(\pm 1|\sigma_k) = (1 \pm \langle\sigma_k\rangle)/2$. Using the properties of the vec operation, we can write this as an explicit function of $\boldsymbol{x}$:

$$\Pr(\pm 1|X, \sigma_k) = \frac{1}{2}\left(1 \pm \frac{\mathrm{Tr}(XX^\dagger\sigma_k)}{\mathrm{Tr}(XX^\dagger)}\right) = \frac{1}{2}\left(1 \pm \frac{\boldsymbol{x}\cdot(\sigma_k \otimes \mathbb{1})\boldsymbol{x}}{\boldsymbol{x}\cdot\boldsymbol{x}}\right) = \Pr(\pm 1|\boldsymbol{x}, \sigma_k). \tag{27}$$

We label the parameter vector within the rank r model M_r by $\boldsymbol{x}_r$ and the associated density matrix given by Eq. (25) ρ_r. Within each model, there is not a one-to-one correspondence between $\boldsymbol{x}_r$ and ρ_r—different vectors will yield the same density matrix. Since we will be interested in obtaining accurate estimates of density matrices as quantified by some norm on the space of density matrices, we will average the models over their ρ_r's rather than their $\boldsymbol{x}_r$'s.

Within each model, then, we have the mean density matrix (recall D is data, C is the context—that is, which Paulis were measured)

$$\hat{\rho}_r(D;C) = \mathbb{E}_{\boldsymbol{x}_r|D;C}\left[\frac{X_r X_r^\dagger}{\mathrm{Tr}(X_r X_r^\dagger)}\right]. \tag{28}$$

Explicitly, the MAE in Eq. (19) is

$$\rho_{\mathrm{MAE}}(D;C) = \sum_r \hat{\rho}_r(D;C)\Pr(r|D;C). \tag{29}$$

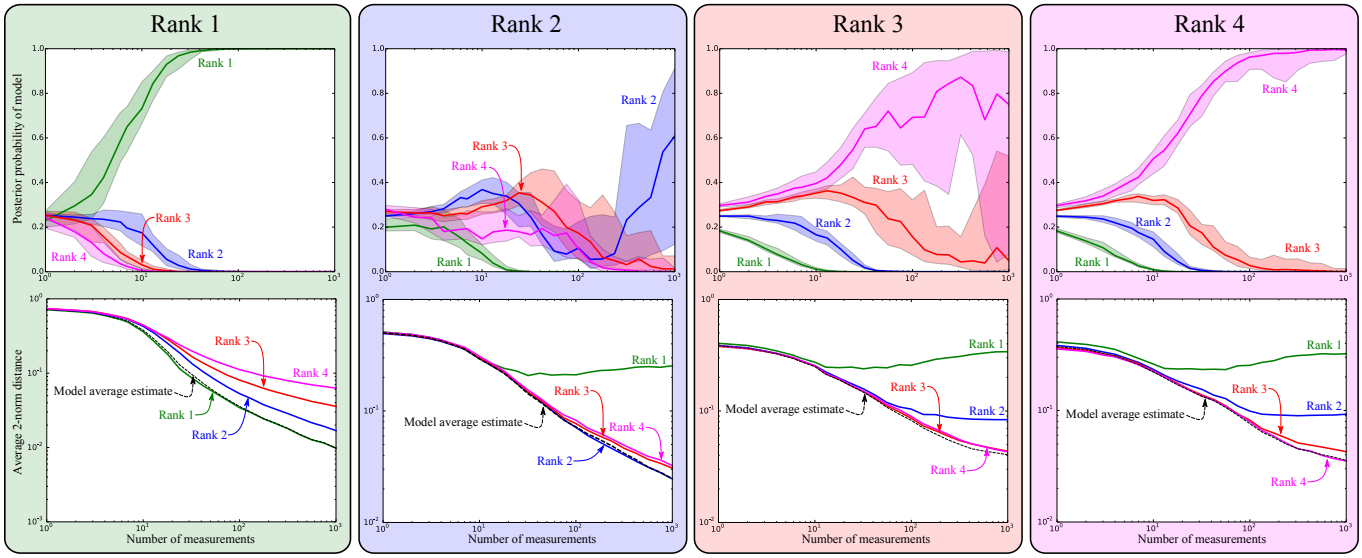

FIG. 1: The performance of Bayesian model selection and the model average estimate for 2 qubits. Each box represents the data for its labeled rank as the "true" model—from left to right, ranks 1 through 4. The lines represent the median of the data and, where present, the shaded areas are the interquartile ranges. Each "measurement" on the horizontal axis corresponds to 100 experiments of a randomly chosen Pauli measurement. In the SMC algorithm, 10^4 particles were used in each model. For each true rank, 100 simulated states were generated and measured.

We assume there is a true density matrix ρ_{t} and we judge each estimate of the true state $\hat{\rho}$ by its spectral distance to ρ_{t}:

$$\|\rho_{\mathrm{t}} - \hat{\rho}\| = \sigma_{\max}(\rho_{\mathrm{t}} - \hat{\rho}), \tag{30}$$

where $\sigma_{\max}$ denotes the largest singular value. This is the norm induced by the usual Euclidean norm on vectors.

The data for 2 qubits is presented in Fig. 1. The important take-away is that the model average estimate does as well, or slightly better than the true model in every case. Also, we see that it is quite easy to identify rank 1 models (pure states) as well as rank 4 models (full rank states) while it seems difficult to identify states of non-extreme rank. Notice that it is extremely difficult to distinguish the rank 3 model from the full rank model when the former defines the underlying truth. However, both models perform well with respect to the error in the estimated parameters and the model average estimate does best, on average.

To further illustrate the difficulty in differentiating high rank states, the probabilities assigned to 3 qubit models is shown in Fig. 2. Again, we see that pure states and full rank states are correctly identified, yet it is difficult to correctly distinguish between rank 7 and rank 8 states, in the same way as it was difficult to distinguish rank 3 and rank 4 states for 2 qubits. We conclude that for the models considered here, it is easiest to correctly identify low rank and full rank states, while it is difficult to correctly identify nearly high rank states.

Notice also that models which are far away—in the sense that the ranks differ by relatively large amounts—are quickly ruled out. In these cases, since the SMC algorithm can be run online (in parallel with the experiment), simulating such models can be stopped to mitigate the computational difficultly in simultaneously simulating many quantum models. Still, tomography is at one extreme in the spectrum of methods for estimating quantum mechanical parameters—it is the *most complete* description of the physical system. At the other extreme is summarizing information from experiments into a single number, such as fidelity.

B. Randomized benchmarking

In the second example, we consider the experimental protocol of randomized benchmarking [59] which has been demonstrated in a variety of experimental settings [61–64] to efficiently characterize noise and quantum channels. The protocol consists of stringing together potentially long sequences of gates which is then undone to determine if the initial state has survived. In [59] the approach was shown to give, in expectation, an exponentially decay $Ap^m + B$, where A and B encode the errors in preparation and measurement, p is the bare survival probability and m is the length of the sequence. In the models we consider, p can be related to the average fidelity over a group of gates, but more specialized protocols exists [24, 65].

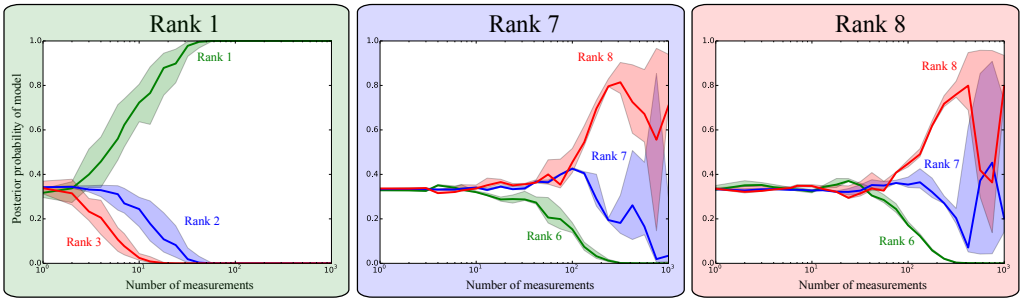

FIG. 2: The performance of Bayesian model selection and the model average estimate for 3 qubits. Each box represents the data for its labeled rank as the "true" model. The lines represent the median of the data and, where present, the shaded areas are the interquartile ranges. Each "measurement" on the horizontal axis corresponds to 100 experiments of a randomly chosen Pauli measurement. In the SMC algorithm, 10^5 particles were used in each model. For each true rank, 10 simulated states were generated and measured.

Typically, p is the only parameter of interest since it is directly related to the average fidelity of the device which is then compared to some threshold. The other parameters are often consider nuisance parameters. In [66] it was shown that the decay can be interpreted as probabilistic model where the binary outcome of each measurement sequence of length m has probability of survival (labeled 0)

$$\Pr(0|\boldsymbol{x}_0; m, M_0) = A_0 p^m + B_0, \tag{31}$$

where the subscripts refer to this as the *zeroth order model*. In [59], a hierarchy of models was introduced because the zeroth order model assumes the errors in the gates within the sequence are independent. By dropping this assumption, a richer set of noise models can be studied [60]. Here we will study the zeroth order model and the *first order model*,

$$\Pr(0|\boldsymbol{x}_1; m, M_1) = A_1 p^m + B_1 + C_1(m-1)(q_1 - p^2)p^{m-2}, \tag{32}$$

where A_1 and B_1 again encode the preparation and measurement errors, C_1 encodes the error on the final gate in the sequence and $q_1 - p^2$ is a measure of the gate dependence in the errors.

For the zeroth order model, we take the prior to be a normal distribution with a mean vector $(p, A_0, B_0) = (0.95, 0.3, 0.5)$ and equal diagonal covariances given by a deviation of $\sigma = 0.01$. Note that the first order model is equal to the zeroth order model when either $C_1 = 0$ or $q_1 = p^2$. In order to not make the model so different that it would be trivial to distinguish them, we look at two priors for the first order model which are close to the zeroth order model. The first is a normal distribution with a mean vector $(p, A_1, B_1, C_1, q_1) = (0.95, 0.3, 0.5, 0.03, 0.95) =: \mu_{\mathrm{I}}$ and equal diagonal covariances given by a deviation of $\sigma = 0.01$ and the second slightly closer with the same covariance matrix but mean vector $(p, A_1, B_1, C_1, q_1) = (0.95, 0.3, 0.5, 0.02, 0.92) =: \mu_{\mathrm{II}}$. Note that the difference between these two distribution in the relative entropy divergence is only 0.050, and so we might expect them to behave the same. Since the both models are close to the zeroth order model, we expect it to be difficult to distinguish them.

In Fig. 3 we simulate the models noting again that, via the priors, they are very close. This intuition is quantified by the fact that the models are hard to distinguish, regardless of which is true. On the left in Fig. 3, we see that parameters in the first order model are so close to those in the zeroth order model that it is irrelevant which is chosen, for the purpose of estimating p. However, what is "close" can be deceiving as we see in the right of Fig. 3. Recall that relative entropy from μ_{I} to μ_{II} is only 0.050 (which explains why they are so difficult to distinguish). In this case, the accuracy of the estimates of the parameter p depend crucially on which model is actually correct. In such cases, the model average estimate can be seen as providing a more conservative estimate of the average gate fidelity by hedging what is at best a 50/50 guess on which model is correct.

V. CONCLUSION AND DISCUSSION

We have introduced a Bayesian model averaging approach to estimating parameters in quantum mechanical models describing data. In the examples considered, the model average estimate performance as well as the unknown true model in most cases. In situations where models are difficult to distinguish, the model average estimate can slightly *outperform* the true model.

For the quantum state estimation example (Sec. IV A) we considered models of differing rank of the density matrix. Ranks which differ by large amounts from the true rank are rapidly ruled out—that is, the probability assigned to

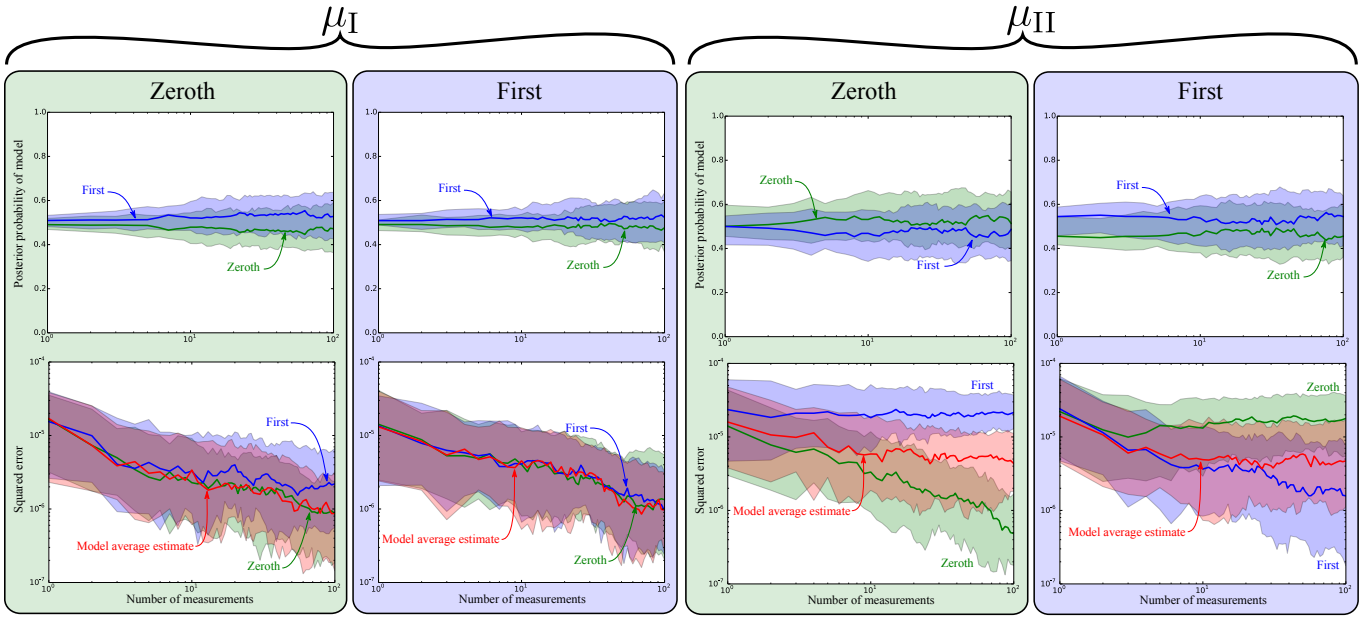

FIG. 3: The performance of Bayesian model selection and the model average estimate for the survival probability in randomized benchmarking experiments. Each box represents the data for its label as the "true" model. The lines represent the median of the data and, where present, the shaded areas are the interquartile ranges. Each "measurement" on the horizontal axis corresponds to a randomized benchmarking experiment with sequences lengths $\{10, 30, 50, \ldots, 200\}$ and 1000 repetitions per sequence length. In the SMC algorithm, 10^3 particles were used in for each model.

them quickly approaches zero. Thus, they contribute nothing to the model average estimate. On the other hand, ranks which are close to the true rank—especially when the true rank is high—are not so easily distinguished. This means that first selecting a rank and then performing estimation within that model can lead to overconfident estimates of the state. By averaging, we can allow those estimates to only contribute with the relative probability that we deem them to be true.

The mechanism for why higher rank states are hard to distinguish is yet unclear. Although the SMC algorithm and the implementation used here have been extensively studied for a wide range of problems, it is possible that state tomography is an outlier. That is, some major modification to modeling or the algorithm itself may be required to obtain the best performance. This resolution would be less interesting than a physical explanation, such as Pauli measurement tomography (which we have used in the example here) is not the optimal scheme to distinguish rank. These questions are left for future work.

In the example of randomized benchmarking (Sec. IV B), we explored a situation where the presence of higher order perturbations were difficult to detect. The Bayesian model selection approach accurately predicts this by assigning roughly 50/50 probability assignments to the models. Surprisingly, the "closeness" of the models as measured by our ability to distinguish them does not translate into our ability to accurately estimate parameters common to each. In some cases we can do no more than guess which is correct, yet guessing the wrong model may have disastrous consequences in our ability to accurately infer the parameters of interest. Again, the model average estimate mitigates the risk of improperly "guessing" which model is correct.

We have noted in the introduction the numerous approaches to estimation within quantum theory. This should urge one to ask, are all of these distinct approaches necessary—is there not some unified approach? Yes! The Bayesian framework outlined here is remarkably powerful in its generality. We note that Bayesian ideas have already been put to good use in quantum information theoretic and foundational problems [69–77] as well as for tomographic and parameter estimation problems [26, 78–85] and experimental design [86, 87]. Importantly, the Bayesian algorithm can provide solutions to these problems *online*, while the experiment is running, with the same software tools [56].

Acknowledgments

The author thanks Chris Granade and Robin Blume-Kohout for helpful discussions. This work was supported in part by National Science Foundation Grant No. PHY-1212445 and by the Canadian Government through the NSERC

PDF program.

[1] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, *Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms*, Journal of Magnetic Resonance **172**, 296 (2005).

[2] C. Brif, R. Chakrabarti and H. Rabitz, *Control of quantum phenomena: past, present and future*, New Journal of Physics **12**, 075008 (2010).

[3] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*. Cambridge University Press (2010).

[4] C. M. Caves, *Quantum-mechanical noise in an interferometer*, Physical Review D **23**, 1693 (1981).

[5] J. Aasi, J. Abadie, B. Abbott, R. Abbott, T. Abbott, M. Abernathy, C. Adams, T. Adams, P. Addesso, R. Adhikari, et al., *Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light*, Nature Photonics **7**, 613 (2013).

[6] M. Paris and J. Rehacek, *Quantum state estimation*. Springer, New York (2004).

[7] C. W. Helstrom, *Quantum detection and estimation theory*, Academic Press (1976).

[8] M. Hayashi, *Comparison Between the Cramer-Rao and the Mini-max Approaches in Quantum Channel Estimation*, Communications in Mathematical Physics **304**, 689 (2011).

[9] M. D. Lang and C. M. Caves, *Optimal Quantum-Enhanced Interferometry Using a Laser Power Source*, Physical Review Letters **111**, 173601 (2013).

[10] V. Giovannetti, S. Lloyd and L. Maccone, *Quantum metrology*, Physical Review Letters **96**, 010401 (2006).

[11] O. Barndorff-Nielsen and R. Gill, *Fisher information in quantum statistics*, Journal of Physics A: Mathematical and General **33**, 4481 (2000).

[12] A. Fujiwara and H. Imai, *A fibre bundle over manifolds of quantum channels and its application to quantum statistics*, Journal of Physics A: Mathematical and Theorical **41**, 255304 (2008).

[13] S. Knysh, V. N. Smelyanskiy and G. A. Durkin, *Scaling laws for precision in quantum interferometry and the bifurcation landscape of the optimal state*, Physical Review A **83**, 021804(R) (2011).

[14] B. M. Escher, R. L. de Matos Filho and L. Davidovich, *General framework for estimating the ultimate precision limit in noisy quantum-enhanced metrology*, Nature Physics **7**, 406 (2011).

[15] S. Boixo, S. T. Flammia, C. M. Caves, and J. M. Geremia, *Generalized Limits for Single-Parameter Quantum Estimation*, Physical Review Letters **98**, 090401 (2007).

[16] M. Tsang, H. M. Wiseman, and C. M. Caves, *Fundamental Quantum Limit to Waveform Estimation*, Physical Review Letters **106**, 090401 (2011).

[17] C. Ferrie, *The best Fisher is upstream: data processing inequalities for quantum metrology*, arXiv:1404.3225 (2014).

[18] B. L. Higgins, D. W. Berry, S. D. Bartlett, H. M. Wiseman, and G. J. Pryde, *Entanglement-free Heisenberg-limited phase estimation*, Nature **450**, 393 (2007).

[19] A. Hentschel and B. C. Sanders, *Efficient Algorithm for Optimizing Adaptive Quantum Metrology Processes*, Physical Review Letters **107**, 233601 (2011).

[20] D. Mahler, L. A. Rozema, A. Darabi, C. Ferrie, R. Blume-Kohout, and A. Steinberg, *Adaptive quantum state tomography improves accuracy quadratically*, Physical Review Letters **111**, 183601 (2013).

[21] C. Ferrie and C. E. Granade, *Likelihood-Free Methods for Quantum Parameter Estimation*, Physical Review Letters **112**, 130402 (2014).

[22] S. T. Flammia and Y.-K. Liu, *Direct fidelity estimation from few Pauli measurements*, Physical Review Letters **106**, 230501 (2011).

[23] M. P. da Silva, O. Landon-Cardinal, and D. Poulin, *Practical characterization of quantum devices without tomography*, Physical Review Letters **107**, 210404 (2011).

[24] E. Magesan, J. M. Gambetta, B. R. Johnson, C. A. Ryan, J. M. Chow, S. T. Merkel, M. P. da Silva, G. A. Keefe, M. B. Rothwell, T. A. Ohki, et al., *Efficient Measurement of Quantum Gate Error by Interleaved Randomized Benchmarking*, Physical Review Letters **109** (2012).

[25] Z. Hradil, *Quantum-state estimation*, Physical Review A **55**, R1561 (1997).

[26] R. Blume-Kohout, *Optimal, reliable estimation of quantum states*, New Journal of Physics **12**, 043034 (2010).

[27] R. Blume-Kohout, *Hedged maximum likelihood quantum state estimation*, Physical Review Letters **105**, 200504 (2010).

[28] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, *Quantum state tomography via compressed sensing*, Physical Review Letters **105**, 150401 (2010).

[29] D. Gottesman, *Identifying Stabilizer States*, Quantum Estimation: Theory and Practice, Perimeter Institute for Theoretical Physics (2008).

[30] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, *Efficient quantum state tomography*, Nature Communications **1**, 149 (2010).

[31] G. Tóth, W. Wieczorek, D. Gross, R. Krischek, C. Schwemmer, and H. Weinfurter, *Permutationally Invariant Quantum Tomography*, Physical Review Letters **105**, 250403 (2010).

[32] O. Landon-Cardinal and D. Poulin, *Practical learning method for multi-scale entangled states*, arXiv:1204.0792 (2012).

[33] J. Řeháček, D. Mogilevtsev, and Z. Hradil, *Tomography for quantum diagnostics*, New Journal of Physics **10**, 043022 (2008).

[34] M. Christandl and R. Renner, *Reliable Quantum State Tomography*, Physical Review Letters **109**, 120403 (2012).
[35] J. M. Arrazola, O. Gittsovich, J. M. Donohue, J. Lavoie, K. J. Resch and N. Lütkenhaus, *Reliable entanglement verification*, Physical Review A **87** 062331 (2013).
[36] R. Blume-Kohout, *Robust error bars for quantum tomography*, arXiv:1202.5270 (2012).
[37] J. Shang, H. K. Ng, A. Sehrawat, X. Li, and B.-G. Englert, *Optimal error regions for quantum state estimation*, New Journal of Physics **15** 123026 (2013).
[38] C. Ferrie, *High posterior density ellipsoids of quantum states*, New Journal of Physics **16**, 023006 (2014).
[39] P. Lougovski and S. J. van Enk, *Characterizing entanglement sources*, Physical Review A **80**, 052324 (2009).
[40] J. O. S. Yin and S. J. van Enk, *Information criteria for efficient quantum state estimation*, Physical Review A **83**, 062110 (2011).
[41] L. Schwarz and S. J. van Enk, *Error models in quantum computation: An application of model selection*, Physical Review A **88**, 032318 (2013).
[42] S. J. van Enk and R. Blume-Kohout, *When quantum tomography goes wrong: drift of quantum sources and other errors*, New Journal of Physics **15**, 025024 (2013).
[43] K. Usami, Y. Nambu, Y. Tsuda, K. Matsumoto and K. Nakamura, *Accuracy of quantum-state estimation utilizing Akaikes information criterion*, Physical Review A **68**, 022314 (2003).
[44] M. Guta, T. Kypraios and I. Dryden, *Rank-based model selection for multiple ions quantum tomography*, New Journal of Physics **14**, 105002 (2012).
[45] J. A. Hoeting, D. Madigan, A. E. Raftery and C. T. Volinsky, *Bayesian model averaging: a tutorial*, Statistical Science **14**, 382 (1999).
[46] K. P. Burnham and D. R. Anderson, *Model selection and multimodel inference: a practical information-theoretic approach*, Springer (2002).
[47] H. Jeffreys, *The Theory of Probability*, Oxford University Press (1998).
[48] R. E. Kass and A. E. Raftery, *Bayes factors*, Journal of the American Statistical Association **90**, 773 (1995).
[49] N. Wiebe, C. Granade, C. Ferrie, and D. G. Cory, *Quantum Hamiltonian learning using imperfect quantum resources*, Physical Review A **89**, 042314 (2014).
[50] J. Combes, C. Ferrie, C. Cesare, M. Tiersch, G. J. Milburn, H. J. Briegel and C. M. Caves, *In-situ characterization of quantum devices with error correction*, arXiv:1405.5656 (2014).
[51] A. Doucet, S. Godsill and C. Andrieu, *On sequential Monte Carlo sampling methods for Bayesian filtering*, Statistics and Computing **10**, 197 (2000).
[52] C. E. Granade, C. Ferrie, N. Wiebe, and D. G. Cory, *Robust online Hamiltonian learning*, New Journal of Physics **14**, 103013 (2012).
[53] N. Wiebe, C. Granade, C. Ferrie, and D. G. Cory, *Hamiltonian Learning and Certification Using Quantum Resources*, arXiv:1309.0876 (2013).
[54] B. A. Chase and J. M. Geremia, *Single-shot parameter estimation via continuous quantum measurement*, Physical Review A **79**, 022314 (2009).
[55] F. Huszr and N. M. T. Houlsby, *Adaptive Bayesian quantum tomography*, Physical Review A **85**, 052120 (2012).
[56] C. Granade, C. Ferrie, et al., *QInfer: Library for Statistical Inference in Quantum Information* (2012), URL `https://github.com/csferrie/python-qinfer`.
[57] E. Jones, T. Oliphant, P. Peterson, et al., *SciPy: open source scientific tools for python* (2001), URL `http://www.scipy.org/`.
[58] K. Zyczkowski and H.-J. Sommers, *Induced measures in the space of mixed quantum states*, Journal of Physics A: Mathematical and General, **34**, 7111 (2001).
[59] E. Magesan, J. M. Gambetta, and J. Emerson. *Scalable and robust randomized benchmarking of quantum processes*, Physical Review Letters **106**, 180504 (2011).
[60] E. Magesan, J. M. Gambetta, and J. Emerson, *Characterizing quantum gates via randomized benchmarking*, Physical Review A **85**, 042311 (2012).
[61] C. A. Ryan, M. Laforest, and R. Laflamme, *Randomized benchmarking of single- and multi-qubit control in liquid-state NMR quantum information processing*, New Journal of Physics **11**, 013034 (2009).
[62] J. M. Chow, J. M. Gambetta, L. Tornberg, J. Koch, L. S. Bishop, A. A. Houck, B. R. Johnson, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, *Randomized Benchmarking and Process Tomography for Gate Errors in a Solid-State Qubit*, Physical Review Letters **102**, 090502 (2009).
[63] S. Olmschenk, R. Chicireanu, K. D. Nelson, and J. V. Porto, *Randomized benchmarking of atomic qubits in an optical lattice*, New Journal of Physics **12**, 113007 (2010).
[64] K. R. Brown, A. C. Wilson, Y. Colombe, C. Ospelkaus, A. M. Meier, E. Knill, D. Leibfried, and D. J. Wineland, *Single-qubit-gate error below* 10^4 *in a trapped ion*, Physical Review A **84**, 030303 (2011).
[65] O. Moussa, M. P. da Silva, C. A. Ryan and R. Laflamme, *Practical experimental certification of computational quantum gates using a Twirling procedure*, Physical review Letters **109**, 070504 (2012).
[66] C. Granade, C. Ferrie and D. G. Cory, *Accelerated Randomized Benchmarking*, arXiv:1404.5275 (2014).
[67] S. Kimmel, M. P. da Silva, C. A. Ryan, B. R. Johnson, and T. Ohki, *Robust Extraction of Tomographic Information via Randomized Benchmarking*, Physical Review X **4**, 011050 (2014).
[68] J. M. Epstein, A. W. Cross, E. Magesan, and J. M. Gambetta, *Investigating the limits of randomized benchmarking protocols*, arXiv:1308.2928 (2013).
[69] R. Schack, T. A. Brun, and C. M. Caves, *Quantum Bayes rule*, Physical Review A **64**, 014305 (2001).

[70] C. M. Caves, C. A. Fuchs, and R. Schack, *Unknown quantum states: The quantum de Finetti representation*, Journal of Mathematical Physics **43**, 4537 (2002).

[71] C. M. Caves, C. A. Fuchs, and R. Schack, *Quantum probabilities as Bayesian probabilities*, Physical Review A **65**, 022305 (2002).

[72] C. A. Fuchs and R. Schack, *Priors in Quantum Bayesian Inference*, AIP Conference Proceedings **1101**, 255 (2009).

[73] C. A. Fuchs and R. Schack, *Quantum-Bayesian Coherence*, arXiv:0906.2187 (2009).

[74] C. A. Fuchs, *QBism, the Perimeter of Quantum Bayesianism*, arXiv:1003.5209 (2010).

[75] C. A. Fuchs and R. Schack, *Bayesian Conditioning, the Reflection Principle, and Quantum Decoherence*, arXiv:1103.5950 (2011).

[76] M. S. Leifer and R. W. Spekkens, *A Bayesian approach to compatibility, improvement, and pooling of quantum states*, arXiv:1110.1085 (2011).

[77] M. S. Leifer and R. W. Spekkens, *Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference*, Physical Review A **88**, 052130 (2013).

[78] K. Jones, *Principles of quantum inference*, Annals of Physics **207**, 140 (1991).

[79] R. Derka, V. Buzek, G. Adam, and P. L. Knight, *From quantum Bayesian inference to quantum tomography*, Journal of Fine Mechanics and Optics **11-12**, 341 (1996).

[80] V. Bužek, *Reconstruction of Liouvillian superoperators*, Physical Review A **58**, 1723 (1998).

[81] B. Teklu, S. Olivares, and M. G. A. Paris, *Bayesian estimation of one-parameter qubit gates*, Journal of Physics B: Atomic, Molecular and Optical Physics **42**, 035502 (2009).

[82] R. D. Gill, *Conciliation of Bayes and Pointwise Quantum State Estimation: Asymptotic information bounds in quantum statistics*, arXiv:math/0512443 (2005).

[83] A. Sergeevich, A. Chandran, J. Combes, S. Bartlett, and H. Wiseman, *Characterization of a qubit Hamiltonian using adaptive measurements in a fixed basis*, Physical Review A **84** (2011).

[84] D. Brivio, S. Cialdi, S. Vezzoli, B. T. Gebrehiwot, M. G. Genoni, S. Olivares, and M. G. A. Paris, *Experimental estimation of one-parameter qubit gates in the presence of phase diffusion*, Physical Review A **81**, 012305 (2010).

[85] D. Oi and S. Schirmer, *Quantum system characterization with limited resources*, Philosophical Transactions of the Royal Society A 28 **370**, 5386 (2012).

[86] C. Ferrie, C. E. Granade, and D. G. Cory, *Adaptive Hamiltonian estimation using Bayesian experimental design*, AIP Conference Proceedings **1443**, 165 (2012).

[87] C. Ferrie, C. E. Granade, and D. G. Cory, *How to best sample a periodic probability distribution, or on the accuracy of Hamiltonian finding strategies*, Quantum Information Processing **12**, 611 (2013).